

INTEGRATING LOCAL AND GLOBAL NAVIGATION IN UNMANNED GROUND VEHICLES

Juan Pablo Gonzalez*, William Dodson, Robert Dean

General Dynamics Robotic Systems

Westminster, MD

Alberto Lacaze, Leonid Sapronov

Robotics Research

Gaithersburg, MD

ABSTRACT

Hierarchical approaches to autonomous navigation usually divide path planning in two levels: local and global navigation. While these two approaches are complementary and can perform very well, they introduce the additional challenge of integrating them in a way that maximizes their strengths and minimizes their weaknesses. In this paper, we evaluate three different approaches to integrating global and local navigation: route-based navigation, route-based navigation with replanning, and combined navigation using the *Field Cost Interface (FCI)*.

1. INTRODUCTION

Hierarchical approaches to autonomous navigation usually divide path planning in two levels: local and global navigation. Local navigation considers the kinematic constraints of the vehicle, and has a sensor-based, high-resolution, near-field representation of the environment. Global navigation typically neglects the kinematic constraints of the vehicle and uses a lower-resolution but farther-reaching representation of the environment while taking mission considerations into account in order to develop tactical plans over longer distances.

While these two levels are complementary and can perform very well, they introduce the additional challenge of integrating them in a way that maximizes their strengths and minimizes their weaknesses. In this paper, we evaluate three different approaches to integrating global and local navigation: route-based navigation, route-based navigation with replanning, and combined navigation using the *Field Cost Interface (FCI)*.

1.1 Route-based navigation

In route-based navigation the global planner generates an initial route using prior map information and taking mission considerations into account. The local planner

then attempts to follow that route while considering the sensor information collected along the route. We use the Geometric Path Planner (GPP) (Gonzalez et al, 2006) in order to generate global routes that consider tactical mission requirements such as travel time, mobility cost, exposure risk and coverage. The local planner is an ego-graph-based planner (Lacaze et al, 1998) that considers the kinematic and non-holonomic constraints of the vehicle. See Fig 1 for an example of a global route generated by the GPP.

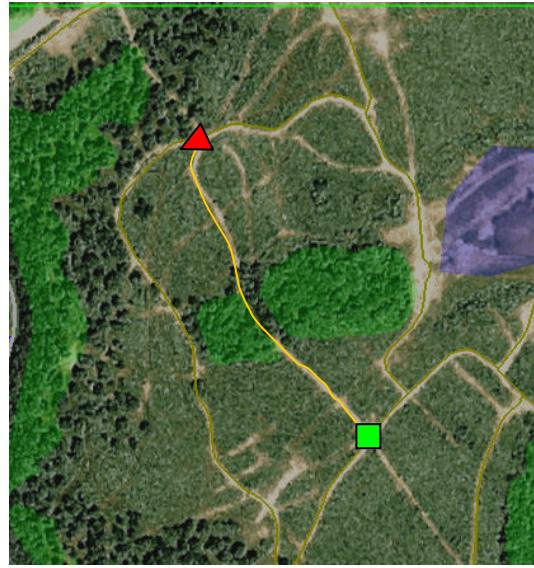


Fig 1. Initial route (orange) planned by global planner.

The main advantage of route-based navigation is that the route that the vehicle will follow is known in advance. However, if this route is not valid, the robot has very limited options to choose a new route. In general, the robot is given a buffer zone around the route and it is allowed to avoid obstacles and find alternate routes around this buffer zone. Fig 2 shows an example of a large blockage that invalidates the global route. In route-based navigation the global route remains the same in spite of such blockages which makes it harder for the local planner to find a viable alternative.

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE DEC 2008	2. REPORT TYPE N/A	3. DATES COVERED -
4. TITLE AND SUBTITLE Integrating Local And Global Navigation In Unmanned Ground Vehicles		
5a. CONTRACT NUMBER		
5b. GRANT NUMBER		
5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)		
5d. PROJECT NUMBER		
5e. TASK NUMBER		
5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) General Dynamics Robotic Systems Westminster, MD		
8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		
10. SPONSOR/MONITOR'S ACRONYM(S)		
11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited		
13. SUPPLEMENTARY NOTES See also ADM002187. Proceedings of the Army Science Conference (26th) Held in Orlando, Florida on 1-4 December 2008, The original document contains color images.		
14. ABSTRACT		
15. SUBJECT TERMS		
16. SECURITY CLASSIFICATION OF:		
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified
17. LIMITATION OF ABSTRACT UU		
18. NUMBER OF PAGES 5		
19a. NAME OF RESPONSIBLE PERSON		

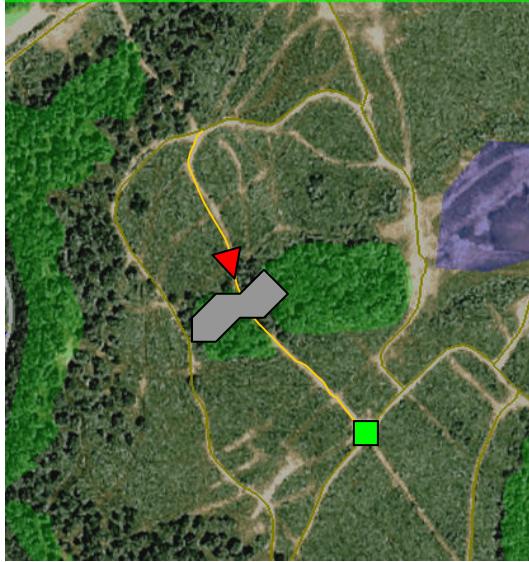


Fig 2. Blockage on route. In route-based navigation the global route does not change in response to sensed data and the robot only uses local navigation to avoid sensed obstacles.

1.2 Route-based navigation with dynamic replanning

In route-based navigation with replanning the global planner generates an initial route as well, but it also incorporates sensor data as the vehicle traverses the route, along with information on the dynamic tactical environment and mission progress, and updates the route periodically. Because the GPP planner uses D* (Stentz, 1995) or Field D* (Ferguson and Stentz, 2006) as its planning algorithms, much of the original search performed by the GPP is reused, enabling for very fast replanning in response to these changes in the environment.

If the robot encounters a blockage or a cul-de-sac, the global planner will be able to find an alternate path based on the combined prior and sensor data available. This route, however does not consider the kinematic constraints of the vehicle and the robot may not be able to follow it. In spite of this, most of the time the local planner is able to smooth the route and produce a valid route for the vehicle to follow. Fig 3 shows an example of the dynamic replanner finding a route to avoid the blockage detected in the sensed data. Notice that the robot would have to first turn around in order to follow the route which would make the proposed route not as desirable as it seems from the global planner's perspective.

1.3 Combined navigation using FCI

When using combined navigation using the FCI the global planner continuously generates a cost field at a radius R from the vehicle, using both prior data and

sensor data. The local planner then attempts to plan paths to each point along this circle, thereby combining the kinematic constraints of the vehicle and the recommendations of the global planner. While planning algorithms used at the global and local level are the same as in the previous approaches, the combination through the FCI provides an interface in which the interactions between the two planners are limited to the boundary of the cost field, in a similar fashion to the approach proposed in (Lacaze, 2002).

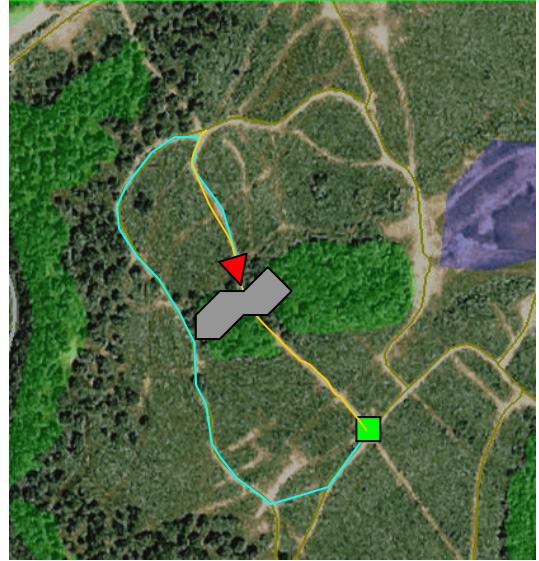


Fig 3. Blockage on route and new route generated automatically by dynamic replanner (cyan)

If there is only one good global route (or a global route that is much better than any alternatives) this approach would produce a similar result to the dynamic replanner, due to the relatively low cost of the single good global route. However, if there are several good global routes with comparable costs, this approach will choose the one that is cheapest from the local planner's point of view, avoiding situations where a route that is non-traversable is chosen.

The main challenge with using FCI is that it requires combining different cost metrics for the global and local planner. In order to combine these costs, the planner first scales both the global and local costs by calculating the average cost assigned to a given section of the route by both planners. This produces a cost metric that has similar scales and that adapts to different terrain types. The planner then combines the scaled costs as follows

$$C_{total} = C'_{local} + k \cdot C'_{global} \quad (1)$$

where C'_{local} and C'_{global} are the scaled local and global path costs, and k is a constant that is determined experimentally. This constant defines the relative weight

of the global costs with respect to the local costs and is the most important parameter for the performance of the algorithm, as it compensates for any systematic differences between the local and global costs.

Fig 4 shows how the FCI evaluates routes for the example from the previous figures. The cyan lines show the global paths being used to generate the field costs, and the yellow line shows the local path chosen after considering both the global and local costs.

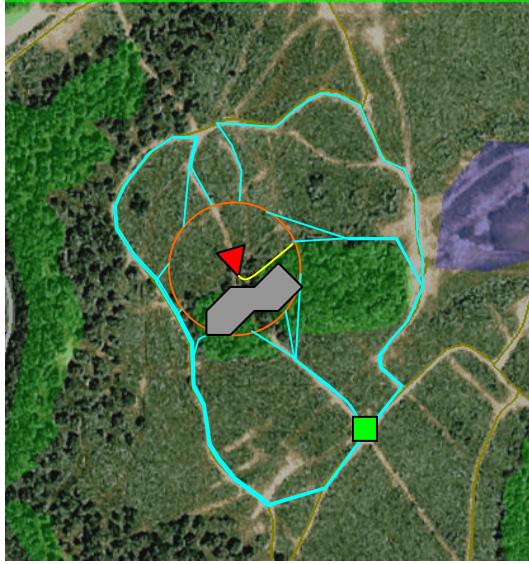


Fig 4. Blockage on route and alternatives passed to the local planner by the FCI (cyan). The new route is selected considering the global paths and the local kinematic constrains of the vehicle.

2. TECHNICAL APPROACH

In order to evaluate the performance of each approach in a controlled manner, a simulated environment was configured such that all the relevant elements to the integration of local and global navigations were properly represented: perception, local planning and global planning. The Robotics Interactive Visualization and Exploitation Technology (RIVET) simulator generated a world scenario, simulated laser returns and simulated the motion of the vehicle through the world. This simulator sent the simulated laser returns to the perception module, which in turn performed terrain classification for the local and global planners. Fig 5 shows the camera view in the RIVET simulator for one of the experiments.

Several experiments were performed in this setup, with different paths, obstacle configurations, densities of obstacles, etc. Fig 6 shows one of the experiments, which illustrates different aspects of the integration between local and global navigation. The path starts in the top right corner, goes through two intermediate waypoints and then finishes near the bottom right corner. The lower left portion of the path is blocked, but this is not known when the path is initially planned.

3. RESULTS

Each one of the three approaches being evaluated was used to execute the test missions. Since the FCI performance depends greatly on the value of k , three different runs were performed for k values of 1.0, 0.1 and 0.01.



Fig 5. RIVET simulator used to evaluate the different approaches to integrating local and global navigation.

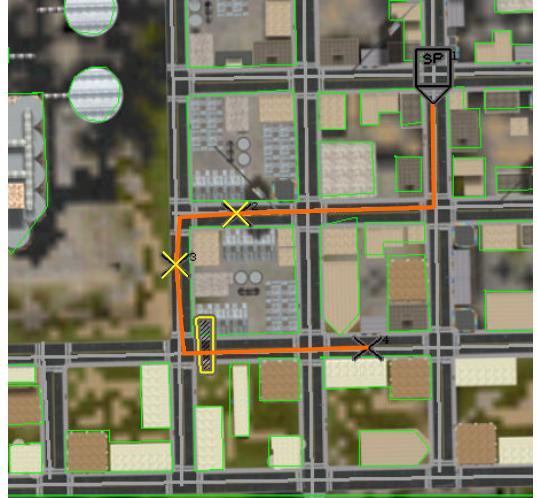


Fig 6. Simulated environment and initial path. The path starts in the top right corner and goes through two intermediate waypoints. The lower left portion of the path is blocked, but this is not known when the path is initially planned.

Fig 7 shows the path executed by the robot using route-based navigation in the mission described by Fig 6. The robot is able to successfully navigate around local obstacles in the beginning of the route, but is unable to find a route around the blockage.

Fig 8 shows the path followed by the robot using route-based navigation with dynamic replanning. This time the robot is able to successfully find a route around the blockage thanks to the global planner and the prior map information. However, while exploring for alternative routes around the blockage, the global routes given to the robot are not consistent with the kinematic constraints of the vehicle and the robot ends up turning around unnecessarily before finding the final route that allows the robot to continue to the goal.

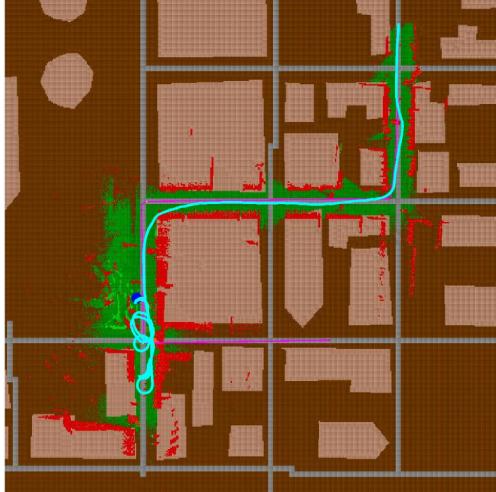


Fig 7. Path executed using route-based navigation. The vehicle is unable to find a route around the blockage using local navigation.



Fig 8. Path executed using route-based navigation with dynamic replanning. The vehicle is able to find a route around the blockage. However, because the route generated by the global planner cannot be followed at times, the vehicle ends up performing one loop near the blockage before continuing on the route.

The following figures show the results of using FCI to plan through the same environment. Three different values of k were used: 1.0, 0.1 and 0.01.

shows the path executed with $k=1.0$. With this value of k the global costs are weighed approximately the same as the local costs. However, because the global part of the path is usually longer, the net effect is that the global aspect of the path is weighed much more than the local aspect. The robot is able to successfully navigate to the goal, but the trajectory tends to oscillate when there are multiple options with similar global costs. As a result, the robot ends up turning around in the bottom left part of the trajectory.



Fig 9. Path executed using route-based navigation with FCI for $k=1.0$. The vehicle is able to find a route around the blockage. However, the trajectory is very sensitive to changes in the global cost as new obstacles are discovered. The robot turns around in the bottom left corner because a cheaper global path is found (with little consideration for the local cost of turning around)

shows the path executed with $k=0.1$. The robot is able to successfully find its path to the goal. The trajectory is very smooth and with little or no oscillation. When the robot approaches the blockage area, it quickly chooses to go around the block, as this is a better alternative considering both global and local costs.

shows the path executed with $k=0.01$. The robot is no longer able to find its path to the goal. Because global costs are weighed so little, the robot ends up performing mostly locally optimal navigation, which makes the robot go almost straight and miss all turns.

summarizes these experiments in terms of total execution time. When $k=0.1$ the FCI is the best performing approach, with an execution time of 310 seconds, followed closely by route-based navigation with replanning at 340 seconds. FCI with $k=1.0$ comes third, at 420 seconds. FCI with $k=0.01$ and route-based navigation without replanning are both unable to find a route to the goal in this scenario.

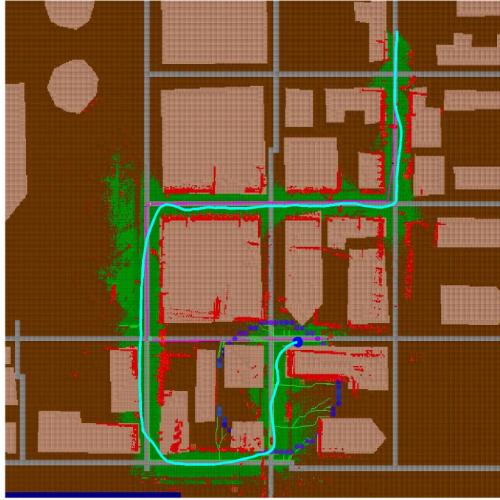


Fig 10. Path executed using route-based navigation with FCI for $k=0.1$. The vehicle is able to find a route around the blockage. The trajectory is smooth and efficient, with a good compromise between global cost and local kinematic constraints.



Fig 11. Path executed using route-based navigation with FCI for $k=0.01$. The vehicle is unable to find a route around the blockage. Because the global costs are weighed so little, the robot ends up performing mostly locally optimal navigation, which makes the robot go almost straight and missing all turns.

Table 1. Execution times of each method

Method	Time (s)
Route-based	N/A
Route-based w/ replanning	340
FCI ($k = 1.0$)	420
FCI ($k = 0.1$)	310
FCI ($k = 0.01$)	N/A

Although these results are for one specific environment, they are representative of our findings in other environments. However, the specific value of k that performs best in a given scenario does change depending on the terrain and other mission-related considerations. In general, FCI is the best performing approach when k is chosen appropriately for a given mission, but can perform very poorly otherwise. Route-based navigation with dynamic replanning does not perform as well as the best FCI runs, but performs well reliably and without depending on any parameters. Route-based navigation without replanning performs well only if original route is valid, but is unable to handle significant deviations from the original route.

4. CONCLUSIONS AND FUTURE WORK

Based on the simulations performed, a *well-tuned* FCI is the best approach for combining local and global navigation. However, if FCI is not well tuned it can perform very poorly. Although route-based navigation didn't perform as well as the well-tuned FCI, it performed consistently well and does not depend on any parameters.

Further experimentation and field validation are still required to better understand when or whether to use each approach. Because the tuning process is still empirical and depending on the specific environment, the preliminary results presented here suggest that route-based navigation with dynamic replanning may be better suited for unknown environments, and that FCI may perform best in known environments that allow for careful tuning of the weights that determine the combination of local and global costs.

REFERENCES

Ferguson, D and Stentz, A., 2006: "Using Interpolation to Improve Path Planning: The Field D* Algorithm". *Journal of Field Robotics*, 2006, 23, 79-101

Gonzalez, J. P.; Nagy, B. and Stentz, A., 2006: "The Geometric Path Planner for Navigating Unmanned Vehicles in Dynamic Environments". *Proceedings ANS 1st Joint Emergency Preparedness and Response and Robotic and Remote Systems*.

Lacaze, A., 2002: "Hierarchical planning algorithms". *Proceedings of SPIE*, SPIE, 2002, 4715, 320

Lacaze, A; Moscovitz, Y.; DeClaris, N. and Murphy, K., 1998: "Path planning for autonomous vehicles driving over rough terrain ". *Intelligent Control (ISIC)*, 1998. Held jointly with IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA), Intelligent Systems and Semiotics (ISAS), Proceedings, 14-17 Sep 1998, 50-55

Stentz, A., 1995: "The Focussed D* Algorithm for Real-Time Replanning". *Proceedings of the International Joint Conference on Artificial Intelligence*.